

The Micro-Level Behavior of Vacancies and Hiring

Steven J. Davis, University of Chicago and NBER

R. Jason Faberman, Bureau of Labor Statistics

John C. Haltiwanger, University of Maryland and NBER

Abstract

Models of labor market search are widely applied to the study of unemployment, worker turnover, wage dispersion and other labor market phenomena. These models afford a central role to the concept of a job vacancy, often treating vacancies as an essential and costly input into the worker recruitment process. Yet, the empirical evidence thus far has been limited to aggregate measures of the vacancy rate. We detail micro-level vacancy behavior using establishment data from the new BLS Job Openings and Labor Turnover Survey. We focus on hires, vacancies and the vacancy *yield* (i.e., the success of a vacancy in generating a hire) and document tremendous heterogeneity in their behavior across industries and at the micro-level. We find the vacancy yield to be countercyclical and nonlinearly increasing in establishment growth. To identify some of the underlying, unobserved factors influencing vacancy posting and worker recruitment, we develop a simple yet flexible stock-flow model that can identify the behavior of the job-filling rate, the flow of vacancies, and the incidence of hires that occur without a vacancy. Our model implies that the job-filling rate is countercyclical and has considerable cross-sectional variation. It also suggests that at least one-third of all hires occur without a vacancy posting. These findings raise a variety of questions about the standard approaches to modeling vacancies and worker recruitment, yet our approach may also aid in resolving several empirical puzzles in the labor search literature.

Keywords: vacancies, hiring, labor market search, establishment microdata

JEL Codes:

1. Introduction

In most models of labor market search, firms post vacancies to attract individuals seeking employment. Once they make contact, the firm and potential employee undergo a matching process. If the returns to each exceed their reservation values, the firm makes an offer, the worker accepts, and a hire is made. A standard assumption in most search models is that firms must post a costly vacancy to attract those searching for a job.¹ This stems from the standard use of a *matching function* to depict the hiring process, which generates a hire using job searchers and vacancies as its “inputs”. In reality, however, other practices may allow firms to recruit workers at little or no cost. Firms may recruit by word-of-mouth, within spot labor markets, or through social networks, and could thus hire without posting any vacancies at all. Others have discussed the “black box” nature of the matching function before (Petrongolo and Pissarides, 2001). While its simplicity is a concern for both of its arguments, job searchers are at least tangible—a person exists regardless of their employment status—and they are essential to creating a hire. Vacancies, on the other hand, are intangible—one only exists empirically if a firm announces it is searching—and it is not necessarily true that one needs a vacancy to produce a hire.

In this paper, we provide a detailed exploration of how and to what extent vacancies lead to hires. The sparse literature on vacancies has thus far only built on the core assumption that the matching process requires a firm to post a vacancy. A greater challenge for empirical work in this area, however, has been the paucity of vacancy data,

¹ These include random search models such as Pissarides (1985) and Mortensen and Pissarides (1994), directed search models such as Moen (1996), and on-the-job search models such as Burdett and Mortensen (1998) and Nagypal (2005). For a complete review of labor market search theory, see Mortensen and Pissarides (1999), Rogerson, Shimer and Wright (2005), and Yashiv (2006).

particularly in the U.S. The pioneering work of Abraham (1983, 1987), and Blanchard and Diamond (1989) creatively used the Help Wanted Index (HWI) as a proxy for vacancies and much of the subsequent research in the U.S. on vacancies builds upon this approach.² The Help Wanted Index is a monthly measure of help-wanted advertisements that appear in U.S. newspapers, produced by the Conference Board. While this series yields sensible empirical patterns (see, e.g., Abraham, 1987; Blanchard and Diamond, 1989; and Shimer, 2005a), its design does not allow a disaggregated approach and hence cannot accommodate a firm-level analysis.

Recently, the Bureau of Labor Statistics (BLS) developed a new survey that directly measures micro-level vacancy postings. The Job Openings and Labor Turnover Survey (JOLTS) has estimates of hires, separations, quits, layoffs and vacancy postings from a monthly survey of establishments. The aggregate JOLTS data have already proven useful in their own right.³ We take the data one step further by exploiting the establishment-level JOLTS data to study the micro-level relationship between vacancies and hires.

We start with a review of the aggregate evidence on vacancies and hires, and in doing so, introduce the concept of the *vacancy yield*, which is the number of hires produced per vacancy posting. To borrow from a production function analogy, one can think of this yield as measuring the efficiency, or average “productivity”, of a vacancy. A standard matching function suggests that the yield is countercyclical, and we find evidence consistent with this notion in both the JOLTS and supporting data.

² Exceptions include Holzer (1994), who studies vacancy rates using a sample of U.S. firms. Coles and Smith (1996) and Yashiv (2000) study vacancy behavior using British and Israeli data, respectively. The latter studies are able to appeal to more structured labor market institutions (e.g., job centers) that are generally not present in the U.S.

³ For example, see Hall (2005a), Shimer (2005a), and Valetta (2005).

We then move to the cross-section and find considerable variation in hires, vacancies and the vacancy yield across broad industry and size categories. At the micro-level, the patterns become more complex, and we identify patterns and variations that much of the labor market search literature has yet to consider. First, we find that in a given month, most establishments post no vacancies and when they do, most post only one. We find that many establishments hire without any reported vacancy, so while most establishments have no employment change, it cannot be the entire explanation. The extent of these patterns varies by industry and establishment size, with high-turnover establishments tending to be those most likely to hire without a vacancy posted. Second, at the establishment-level, we find that hires, vacancies, and the vacancy yield are all increasing functions of employment growth, but that these relations are highly nonlinear. Among contracting establishments, the relationship of each to growth is essentially flat, while among expanding establishments, all three exhibit steep increases with growth. Hires, vacancies, and the vacancy yield are lowest among stable establishments (i.e., no employment change), producing a sharp nonlinearity in the growth relationship. Controlling for establishment-specific differences in the data alters the vacancy-growth relation somewhat, but has essentially no effect on the relations of hires or the vacancy yield to growth.

Even with the microdata, we must deal with hires and vacancy measurement issues. In particular, we must address the fact that hires are measured as a *flow* during the month, while vacancies are measured as a *stock* at the end of the month. To get at the processes underlying our empirical observations (e.g., vacancy durations, vacancy flow rates, and the propensity of hiring without a vacancy posting), we introduce some

structure to our empirical analysis with a simple stock-flow model of vacancy posting and worker recruitment. The model is largely an accounting of hires and vacancy dynamics that depicts the monthly behavior of the vacancy stock and hires flow built up from a flow of daily postings and job filling. When confronted with the data, it yields estimates of the (empirically unobservable) flow of vacancy postings and the job-filling rate (the inverse of vacancy duration). Moreover, the model is flexible enough that we can expand it to allow for heterogeneity across establishments and hires to occur without a posted vacancy.

Our stock-flow model, in both a basic and expanded form, yields several findings. The JOLTS data suggest that the fill rate for vacancies is relatively volatile at monthly frequencies, implying substantial variation in vacancy duration. Supporting data suggest that the fill rate, like the vacancy yield, is countercyclical. Further, our basic model suggests that much of the variation in vacancy stocks hires flows stems from movements in the fill rate, not the vacancy flow rate. Once we expand the model to allow for heterogeneity and hiring without a vacancy, however, movements in hires and the vacancy stock are accounted for about equally by the vacancy flow rate, the fill rate, and hires without a vacancy. Particularly interesting is the fact that our expanded model predicts that at least 36 percent of hires occur without any vacancy posting. Such a large fraction is remarkable because it occurs even after accounting for the intra-month flow of vacancies.

As one might expect, we also find tremendous variation across industries and size classes in the fill rate, the vacancy flow rate, and share of hires without a vacancy. We also find strong relationships of each to employment growth at the establishment level. In

fact, their patterns mimic the increasing, nonlinear relationships observed with the hires, vacancies, and the vacancy yield data. Of the three, vacancy flows are the most responsive to increases in the growth rate.

Overall, this paper documents cyclical patterns and micro-level heterogeneity in vacancies and recruitment heretofore unexplored in the literature. Our concept of the vacancy yield proves useful in exploring their time-series and cross-sectional behavior, and fits well within a search and matching framework. Our stock-flow model is simple yet flexible enough to address much of the heterogeneity and dynamics we observe in the data, and has the potential to provide insight into several puzzles currently confronting the labor search literature, such as the inability of a standard matching model to generate the observed volatility in vacancies noted by Hall (2005b), Shimer (2005a) and Mortensen and Nagypal (2006).

The following section describes the sample of JOLTS data we use. Section 3 presents the aggregate and micro-level evidence on vacancies and hires. Section 4 presents our stock-flow model and its results, while Section 5 concludes.

2. Data

For much of this study, we use a sample of microdata from the Job Openings and Labor Turnover Survey (JOLTS), produced by the BLS. The JOLTS is relatively new, and contains data on vacancies, hires, and separations, reported directly by establishments, with separations broken out into quits, layoffs and discharges, and other separations (e.g., retirements). The JOLTS samples roughly 16,000 establishments each month to produce published estimates for total nonfarm payrolls and major industry sectors. The data begin in December 2000 and are updated monthly.

Our sample continues through January 2005, and includes all observations of establishments with positive employment in two consecutive months. This minimizes the potential spurious effects of outliers and inconsistent data reporters. It also allows us to compare hires in the current month to vacancies posted in the previous month. The resulting sample contains 372,288 observations, which represents 93 percent of the pooled observations used in the estimation of published data and, due to the requirement of continuous reported employment, excludes the December 2000 observations.⁴

As Faberman (2005) discusses, there are differences with the timing of the measurement of worker flows and employment in the JOLTS data. To address these differences, we force consistency between our growth and worker flow estimates by defining the net change in employment for month t as hires (h_t) – total separations (s_t). We then define employment in the previous month as the current month’s employment (e_t) minus this net change. By doing so, we can measure the net growth rate as the net change divided by the average of the current and previous months’ employment, which provides the symmetric growth rate (bounded between -200 and 200 percent) described in Davis, Haltiwanger, and Schuh (1996). Hires and separations rates are similarly defined. We measure the vacancy rate slightly differently, as vacancies (v_t) divided by the sum of vacancies and average employment.

The JOLTS data have a relatively short time-series. Consequently, we supplement our JOLTS time-series evidence with gross flows data from the Consumer Population Survey (CPS) and the Help Wanted Index (which is a proxy for vacancies) from the

⁴ The restriction to continuous establishments has little effect on the aggregate estimates, so selection is not an issue. A broader issue is the fact that, due to its sample nature, JOLTS fails to capture most entry and exit (which is why our restriction has so little an effect). Faberman (2005) details the JOLTS sampling and other measurement issues. Another issue may be the effect of data imputations on our results, and in future versions of this paper, we hope to address this concern.

Conference Board. Our gross flows estimates come from two sources. The first is Shimer (2005b), who produces a gross flows series that accounts for time-aggregation issues. From these flows, we derive a hires rate defined as the sum of the unemployed and those not in the labor force who become employed during the month (i.e., unemployment-to-employment and not in the labor force-to-employment flows) divided by employment. We detrend the data using a low-bandwidth Hodrick-Prescott filter (smoothing parameter = 10^5) to account for changes to the CPS survey during the period, and to make their cyclical patterns comparable to the HWI data, which we also detrend for reasons described below. The second source is Fallick and Fleischman (2004). Their data differ from the Shimer series in that they include job-to-job transitions (i.e., employment-to-employment flows) which we add to the above hires measure. The Shimer series spans 1967-2004, while the Fallick-Fleishman series spans 1994-2004.⁵ Both series use quarterly averages of monthly values.

The HWI data are a monthly index of the number of help-wanted advertisements listed in a sampling of U.S. newspapers. Using such a measure as a proxy for vacancies has obvious drawbacks, not the least of which is the recent substitution from help-wanted advertising to job postings via the internet. The index nature of the data also mean that fine disaggregation, at the level we do with the JOLTS data, is not feasible. Nevertheless, the HWI is the only proxy of vacancies that has a long, high-frequency time-series. To deal with its shortcomings, we detrend the data again using a low-bandwidth HP filter (smoothing parameter = 10^5). We then rescale the deviations from trend so that their

⁵ We thank Rob Shimer, Bruce Fallick and Charles Fleischman for providing these data. The Shimer flows from 1967-1975 use estimates tabulated by Joe Ritter and Hoyt Bleakley. The Fallick-Fleischman series begins in 1994 because this is when the CPS underwent a major overhaul, which allowed one to calculate employment-to-employment flows from the microdata.

mean matches the mean JOLTS vacancy rate over the JOLTS sample period. Finally, we use quarterly intervals of the data, where each observation represents the vacancy stock at the beginning of the quarter. Note that over this period, the JOLTS vacancy measure and our adjusted HWI measure track each other very closely.⁶

3. Aggregate and Micro-Level Evidence

3.A. Aggregate Evidence

The period covered by our JOLTS sample spans the onset of a recession and its recovery. The recession officially lasted from March to November 2001, but employment losses continued through mid-2003. Figure 1 depicts the time-series movements of hires, quits, layoffs and vacancies. Hires and vacancies, as well as quits, dip during the recession and remain low afterwards. The vacancy rate undergoes the largest decline. When employment growth picks up again in mid-2003, hires, vacancies, and quits follow. Layoffs rise during the recession and decline thereafter. They remain relatively flat through mid-2003, then begin another, more gradual, decline.

Figure 2 depicts the behavior of hires and vacancies from the CPS gross flow data and the Help Wanted Index back to 1967. The figure depicts both the Shimer and Fallick-Fleischman hires series. Note that the latter hires are greater in magnitude because they include job-to-job transitions. Note also that HP filtering removes a secular decline in hiring rates observed in other research (Faberman, 2006; Davis et al., 2006). With these caveats in mind, the figure shows that both hires and vacancies rise in booms and drop in recessions, with the latter being much more cyclically volatile.

⁶ Our adjustment approach follows techniques developed by Abraham (1987) and Shimer (2005a). Both authors discuss the measurement issues of the HWI in detail.

Our focus in this paper is on the vacancy yield. Figure 3 depicts the time-series of the aggregate JOLTS vacancy yield, measured as the flow of hires during month t , divided by the stock of vacancies at the end of month $t-1$. The flow versus stock comparison is a major the reason observed yields are greater than one (we discuss other potential reasons below). The series appears countercyclical, though much of the movement in the vacancy yield seems driven by the precipitous drop in the vacancy rate observed in 2001. Figure 4 uses the CPS gross flow data and the HWI to calculate a crude measure of the vacancy yield using each hires rate series. It also depicts the shorter JOLTS series (adjusted to a quarterly measure) for comparison. These data strongly affirm that the vacancy yield is countercyclical. Given a standard model of labor market search, this is exactly what we would expect. To see this, let hires stem from a constant returns to scale matching function that has the stocks of vacancies and unemployed (u) as its arguments: $h = \mu v^{1-\alpha} u^\alpha$, where $\mu > 0$ and $0 < \alpha < 1$. Rearranging, we get

$$(1) \quad \frac{h}{v} = \mu \left(\frac{v}{u} \right)^{-\alpha}.$$

With equation (1), it is straightforward to see that the vacancy yield (h/v) is a decreasing function of labor market tightness (v/u). In the data (using the unemployment rate from the CPS and the Shimer hires rate) the correlation between these two measures (in logs) is -0.87, while the correlation between the (log) vacancy yield and the (log) unemployment rate is 0.63. Using the monthly JOLTS data, the correlations are -0.83 and 0.82, respectively.

Table 1 reports the cross-sectional evidence of hires, separations, vacancies, and the vacancy yield by major industry and establishment size class from our JOLTS

sample. There is sizable variation in all variables across industries and across size classes. Of particular note are the variations across industry and size of the vacancy yield. Industries such as Construction and Resources have yields that are several times larger than those in Health & Education and Government. Similarly, the vacancy yield tends to decrease with establishment size. Is it that certain industries and establishment sizes are more efficient at matching workers to jobs? Perhaps. A more likely explanation, however, is that there are institutional differences across these groups in how they recruit and attract workers. For example, establishments in construction and resources may regularly recruit workers from a select labor pool for repeated short-term work, reducing their need for a formal vacancy posting to attract workers. On the other hand, establishments in education, health and government may have regulatory constraints that require them to undergo a formal search process for any new employee. Such differences have important theoretical implications because they suggest that the standard assumption that firms must post a costly vacancy to attract a worker may be true in some industries (and size classes) more than others.

3.B. The Micro Behavior of Vacancies

To truly understand aggregate vacancies and hiring one must examine their behavior at the micro-level. The JOLTS data are the first timely, representative data source that allows such an examination. Consequently, it is useful to know basic micro-level evidence on the frequency, intensity, and variability of vacancy posting. We present that evidence here.

First, one must realize that at the monthly frequency, vacancies are relatively rare. Table 2 illustrates this point with both unweighted and employment-weighted estimates.

In the average month, only 12 percent of establishments (representing 54 percent of employment) post a vacancy. Figure 5 shows that, employment-weighted, even when establishments do post vacancies, they are often at very low rates and levels.

Employment-weighted, the vacancy rate at the 90th percentile is 6 percent while the *number* of vacancies at this percentile is 58, but unweighted, vacancy rate is 3 percent, while the number of vacancies is just one.

Much of this stems from the fact that only 18 percent of establishments (representing 64 percent of employment) report any hire in a given month, diminishing their need for a vacancy. Nevertheless, this cannot be the entire story. For instance, Table 2 shows that nearly 1 in 5 vacancies do not lead to a hire the following month, suggesting that the search process takes some time. Yet, 42 percent of hires occur at establishments where there was no vacancy posted going into the month. This implies that the standard approaches to studying hiring and vacancy behavior fails to capture some aspect of recruitment. One facet of recruiting patterns is their variation across industries and establishment size. Table 2 shows that there are considerable differences in the frequency of hiring and vacancy posting across both industries and size classes. Perhaps counterintuitively, industries with the greatest worker turnover (i.e., highest reported hires and separations) also have the highest shares of observations with no reported vacancies. Consequently, these industries have the highest shares of hires without a previously posted vacancy. When establishments in these industries do post a vacancy, however, they are the most likely to have that posting remain unfilled after a month. We explore possible explanations of these patterns in Section 4 below.

3.C. Hires, Vacancies, and Micro-Movements in Net Growth

We next explore the establishment-level relationships of hires and vacancies to employment growth. It is a stylized fact that there is a wide distribution of growth rates at the establishment level at any point in time (Davis, Haltiwanger, and Schuh, 1996). In addition, labor market search theories suggest that the extent of an establishment's employment change is a signal of the intensity of an idiosyncratic shock. As such, examining the relation of hiring and vacancy posting to employment growth can provide insight on how their behavior varies with the extent of such shocks.

Using our pooled sample of JOLTS microdata, we estimate weighted-mean values of the hires rate, vacancy rate and vacancy yield for growth rate intervals that increase with the magnitude of the growth rate. The intervals are relatively fine (0.1 percent) close to zero and increase to 5 percent intervals near the extremes. The infrequent occurrences of large changes coupled with the relatively small size of the JOLTS sample necessitate the non-uniform interval spacing. We take a semi-parametric approach to estimating the mean values by regressing the variable of interest on a set of dummies for each growth rate interval. This allows us to estimate the vacancy and hires relations to growth while controlling for establishment fixed effects.

Figures 6 and 7 illustrate our results for the hiring rate and vacancy rate, respectively.⁷ Both rates increase with growth, though both relations are nonmonotonic. The hires relation must satisfy some portion of an adding-up constraint, since net growth is the difference between hires and separations. Consequently, the minimum for the hires rate is the horizontal axis for non-positive growth and the 45-degree line for positive growth. Hiring lies above the minimum for all growth rates. Rates hover around 3 percent

⁷ For all figures that depict estimates as a function of net growth, we focus on the -30 to 30 percent range, as greater magnitudes have a large decline in the number of observations used in estimation, and consequently a large decline in statistical precision.

of employment for contracting establishments then decline as one approaches zero. Establishments with no net employment changes have an average hires rate of 1.1 percent. Hiring at expanding establishments increases proportionally with growth, and lies several percentage points above the 45-degree line for all values. Interestingly enough, inclusion of establishment fixed effects does little to alter the observed pattern. Vacancies mostly follow the same pattern, with rates at contracting establishments generally averaging 2 percent regardless of the magnitude of the contraction. Vacancy rates increase with growth, but at a much slower rate than hires—establishments that grow by 30 percent have vacancy rates of just 5.6 percent. The most notable contrast with hires, however, is the relatively sharp discontinuity right around zero growth. Establishments with very small contractions average vacancy rates of 2.0 percent, while establishments with very small expansions average vacancy rates of 2.5 percent. Establishments with zero growth, though, have average vacancy rates of just 1.3 percent, but note that this group includes both idle establishments and establishments whose separations offset hires. When we control for establishment effects, much of the nonlinearities in the vacancy-growth relation disappear.

Figure 8 presents the employment-weighted probability of the discrete event that an establishment posts any vacancy as a function of the growth rate. The relationship is highly nonlinear, with establishments with small employment changes being the most likely to post a vacancy. The probability of a posting decreases sharply in the magnitude of the change, though expanding establishments have a considerably higher probability of posting than contracting establishments. While establishments with very small changes have a probability of a vacancy posting near 80 percent, those with no changes have a

probability of posting of only 22.7 percent. When we control for establishment effects, nearly all the nonlinearities disappear, though the probability remains increasing in growth and a much smaller discontinuity for zero-growth establishments still exists. This is consistent with the notion that different types of firms persistently use vacancies in different ways.

In Figure 9, we present the vacancy yield as a function of growth measured two ways. The first is total hires divided by total vacancies within each growth rate interval; this is similar to dividing the hires function in Figure 6 by the vacancy function in Figure 7, and is depicted in the upper panel.⁸ The second is the micro-level number of hires per vacancy averaged across all establishments within each interval that have a vacancy posted. Both measures show a very similar picture. Among contracting establishments, vacancy yields are constant at about one hire per vacancy. There is a discontinuity for zero-growth establishments, with a slight spike upwards for the first measure and a slight spike downwards for the second measure. The former stems from the sharp drop in vacancies posted in Figure 7. The latter reflects only the yield for establishments who had a posted vacancy at the end of the previous month, so it suggests either that stable establishments tend to be less successful when posting vacancies or the more tautological conclusion that establishments that do not fill their vacancy postings by definition do not grow. Among expanding establishments, both measures increase considerably with the growth rate, with expansions in the 25-30 percent range having over five hires per vacancy. Interestingly, even though establishment fixed effects affect the relation between the vacancy rate and the probability of posting a vacancy, they have very little effect on the relation of the vacancy yield to establishment growth.

⁸ It is not identical to this approach because the hires and vacancy rates have different denominators.

One natural consequence of the hires and vacancy measurement in JOLTS is the comparison of a flow to a stock. As a result, many hires likely occur from vacancies that are posted and filled between the monthly point-in-time measurements of the vacancy stock. Nevertheless, it is unclear why or even whether these intra-month vacancy flows would vary with establishment growth. Other aspects of hiring may lead to the same empirical relation. For example, the flow of vacancies may be independent of growth but vacancy durations may decrease with growth, leading to the observed pattern. Further, it may be that the propensity to attract workers without posting a vacancy or the ability to attract multiple hires per vacancy increases with the growth rate. Regardless of the true underlying behavior, the main point is that the relationship between the vacancy yield and growth cannot be a data anomaly; at least one of the above aspects of vacancies must vary with establishment growth, which has considerable implications for how economists envision and model the labor market search and matching process.

4. Modeling Hiring and Vacancy Dynamics

4.A. Model Overview

We now present a simple stock-flow model of vacancy and hiring flows. The model is designed to pin down key parameters, unobservable in the data, that describe the search and recruitment process, while addressing the inherent time-aggregation issues of comparing stock and flow data. Namely, we seek to identify the average daily fill rate of posted vacancies (denoted by f) and the average daily vacancy flow (denoted by θ). The former is the inverse of the vacancy duration rate. The latter provides a flow measure of vacancies, which allows a flow-versus-flow comparison of hires and vacancies. Later, we extend the model to allow hiring without a vacancy posting.

4.B.1. Basic Model

Let $h_{s,t}$ denote the number of hires on day s during month t , and $v_{s,t}$ denote the number of vacancies on day s during month t . We assume an average daily fill rate (f_t) and vacancy flow (θ_t) for a month consisting of τ days. Note that these parameters are constant over the course of any given month but can vary across months. Hires are simply the share of the vacancy stock from the previous day that is subsequently filled:

$$(2) \quad h_{s,t} = f_t v_{s-1,t}.$$

The stock of vacancies evolves along three dimensions. First, the flow of new vacancies increases the stock. Second, the number of hires during that day depletes the stock. Finally, an exogenous number of vacancies that close without ever being filled also deplete the stock. We denote this last variable by δ_t , and again assume a constant value over the month. The daily equation of motion for the vacancy stock is then

$$v_{s,t} = (1 - \delta_t) v_{s-1,t} + \theta_t - h_{s,t},$$

and substituting in (2) we get

$$(3) \quad v_{s,t} = (1 - f_t - \delta_t) v_{s-1,t} + \theta_t.$$

Next, we need to sum up equations (2) and (3) into monthly measures, as this is what we observe in the data. For vacancies, we would like to relate their stock at the end of month $t-1$, v_{t-1} to their stock at the end of the following month, v_t , τ days later. One can add up equation (3) over τ days and substitute back for $v_{s-1,t}$ to get the desired equation

$$(4) \quad v_t = (1 - f_t - \delta_t)^\tau v_{t-1} + \theta_t \sum_{s=1}^{\tau} (1 - f_t - \delta_t)^{s-1}.$$

The first term on the right depicts the original stock after depletion by hires and closings.

The second term represents the total monthly flow of vacancies, similarly depleted.

Observed hires are a flow measure. As such, we wish to add up daily the daily equation

for hires, so that so that the monthly flow is $H_t = \sum_{s=1}^{\tau} h_{s,t}$. Substituting (3) into (2),

and (2) into the monthly sum, and then substituting back for $v_{s-1,t}$ to the beginning of the month yields the following

$$(5) \quad H_t = f_t v_{t-1} \sum_{s=1}^{\tau} (1 - f_t - \delta_t)^{s-1} + f_t \theta_t \sum_{s=1}^{\tau} (\tau - s)(1 - f_t - \delta_t)^{s-1}.$$

The first term on the right represents hires from the original stock, while the second term represents hires from the total monthly flows. Given an exogenous δ_t , we have two parameters to identify: f_t and θ_t . Equations (4) and (5) give us a two-equation system to exactly identify these parameters.

4.B.2. Estimation Approach

We estimate (4) and (5) using the aggregate hires and vacancy estimates constructed from our JOLTS sample, seasonally adjusted. As a robustness check (that also has the benefit of a longer time series), we also present the results using the gross flow and HWI data. We let H_t be hires during month t , v_t be the vacancies posted at the end of month t and v_{t-1} be the vacancies posted at the end of month $t-1$. For simplicity, we assume all months have $\tau = 26$ working days (the average days per month less Sundays and major holidays). We let δ_t equal L_t/τ , where L_t is the layoff rate for month t . This assumption states that vacancies close without being filled at a rate proportional to the daily layoff rate. This is analogous to assumptions in the labor search literature that set an exogenous job separation rate equal to the layoff rate. We solve the system numerically for each month to obtain estimates of f_t and θ_t , which provides us with a time-series of

each parameter. We can calculate the average vacancy duration (in days) as $1/f_t$ and the monthly flow rate of vacancies as $\tau \cdot \theta_t$.

4.B.3 Results for Basic Model

We begin with the time-series results of our basic two-parameter stock-flow model. To make the results more easily comparable to the results reported in section 3, we scale our flow estimates (which are in levels) by employment in month t , e_t . The scaling has no effect on estimates of f_t , but it allows one to interpret θ_t as an average daily flow *rate* for vacancies.

Figure 10 shows the movements of the monthly vacancy flow rate (measured as $\tau \cdot \theta_t$) and the daily fill rate (f_t), as well as the beginning stock of vacancies, v_{t-1} (measured directly from the JOLTS data). The top row of Table 3 reports that the flow of vacancies average 3.4 percent of employment (compared to the stock's average of 2.4 percent), while the fill rate averages 5.5 percent of the previous day's stock. In the data, the stock of vacancies exhibit much greater cyclical movement than hires (see Figure 1). The results from the basic model suggest that the flow of vacancies is much less volatile than their stock, and that movements in the daily job-filling rate seem to account for much of the observed cyclical movement in the vacancy stock. We explore the contribution of each to the movements in hires and vacancies in counterfactual exercises below. During the 2001 recession, the job-filling rate rises from its low of 4.4 percent to a peak of 6.2 percent in mid-2002, and remains relatively high through the remainder of the period. This coincides with the relatively sharp, persistent decline in the vacancy stock, and represents a decrease in average vacancy duration from 23 days to 16 days.

Both as a robustness check and to better gauge the cyclical movements in these parameters, we re-estimate the model using our adjusted CPS gross flows from Shimer (2005b) and adjusted Help Wanted Index data.⁹ The results are in Figure 11. For the longer time series, the striking pattern is the substantial variation in the job-filling rate, which increases considerably around cyclical downturns. Moreover, the variation in the job-filling rate is much greater than the variation in vacancy flows (the coefficient of variation for the job-filling rate is 0.56, compared to 0.05 for the vacancy flow rate). In terms of the overlap period with JOLTS, the patterns are also broadly similar, though there are some subtle but potentially important differences in the variation in the fill rates during the overlap period. With the JOLTS data, the fill rate rises substantially during the recession but remains high throughout the post 2001 period. For the HWI data, the fill rate also rises during the recession but peaks during the recession. Moreover, the fill rate using HWI begins to fall noticeably in 2004, while the JOLTS fill rate remains high during this period—Figure 4 illustrates similar patterns across the data with the vacancy yield. This implies that, depending on the survey studied, cyclical increases in the job-filling rate have may high or low persistence.

In section 3, we documented considerable variation in the use and yield of vacancies across industries and establishment size. The remainder of Table 3 presents estimates for the basic stock-flow model for these categories. Again, there is considerable variation across industries and size. The model suggests that Resources, Construction and Retail Trade, i.e., the industries with the highest worker turnover and lowest incidence of vacancies, have the highest vacancy flow and job-filling rates. Industries with the lowest

⁹ For the CPS data, we only use 1976 forward because of limitations to the availability of separations data. We need the latter to get an estimate of δ_t , and specifically use the employment-to-unemployment flow, divided by employment, and divided again by τ .

flows and highest use of vacancies, Government and Education and Health, have the lowest rates of both. Across size classes, the vacancy flow and job filling-rates both tend to decrease with establishment size.

4.C.1 Extending the Basic Model: Allowing for Unposted Vacancies

Equations (4) and (5) tacitly assume that all hires must come from a posted vacancy. In reality, hires often occur through recruitment by word-of-mouth, social networking, or other informal methods. Furthermore, our evidence above suggests that such hiring practices may play a significant role in the observed increase in the vacancy yield with establishment growth. Consequently, one extension of our model we explore allows establishments to hire from what we call “unposted” vacancies. We define η_t as the average daily flow of hires from unposted vacancies during month t .¹⁰ The addition of η_t makes the basic model, in its current specification, underidentified. Further, the notion of permitting hires from unposted vacancies represents a significant departure from the standard assumptions of the search and matching literature, so there is little guidance on how to specify such hiring. Accordingly, we discuss a range of alternatives for these hires within our extended model specification.

The additional information we bring to bear follows from empirical findings. Recall that most establishments post no vacancies and that these establishments account for 42 percent of all hires. Consequently, we introduce heterogeneity into our model along this dimension by assuming that establishments that begin the month with no

¹⁰ Note that it is not feasible with the information available to identify an unposted vacancy flow rate and an unposted vacancy fill rate separately so we identify only their product.

vacancies to have different vacancy flow rates (θ_t) during the month than those who begin with at least one vacancy posted.¹¹

In particular, let θ_t^P be the average daily posted vacancy flow for those who begin the month with at least one posted vacancy, and let θ_t^0 be the flow for those who begin the month with no posted vacancy. Let H_t^P and H_t^0 be the cumulative flows of hires during month t for establishments with $v_{t-1} > 0$ and $v_{t-1} = 0$, respectively. Finally, let v_t^P and v_t^0 be the stock of posted vacancies at the end of month t for establishments with $v_{t-1} > 0$ and $v_{t-1} = 0$, respectively. Then, for establishments with a positive stock of posted vacancies at the end of month $t-1$, equations (4) and (5) become

$$(4a) \quad v_t^P = (1 - f_t - \delta_t)^\tau v_{t-1}^P + \theta_t^P \sum_{s=1}^{\tau} (1 - f_t - \delta_t)^{s-1}, \text{ and}$$

$$(5a) \quad H_t^P = f_t v_{t-1}^P \sum_{s=1}^{\tau} (1 - f_t - \delta_t)^{s-1} + f_t \theta_t^P \sum_{s=1}^{\tau} (\tau - s)(1 - f_t - \delta_t)^{s-1} + \tau(1 - \rho_t^0)\eta_t$$

where ρ_t^0 denotes the share of hires from unposted vacancies that occur at establishments with $v_{t-1} = 0$. For establishments with no posted vacancy at the end of $t-1$, equations (4) and (5) simplify to

$$(6) \quad v_t^0 = \theta_t^0 \sum_{s=1}^{\tau} (1 - f_t - \delta_t)^{s-1}, \text{ and}$$

$$(7) \quad H_t^0 = f_t \theta_t^0 \sum_{s=1}^{\tau} (\tau - s)(1 - f_t - \delta_t)^{s-1} + \tau \rho_t^0 \eta_t.$$

We can aggregate equations (4a) and (6), and (5a) and (7), respectively, to yield:

¹¹Note that we could instead have introduced heterogeneity into the model via differing fill rates (f) or differing hires rates from unposted vacancies (η). We experimented with such models and found (in results not reported here) that the nonlinearities in the model and the way in which these parameters enter the model make it difficult to solve such specifications numerically. An alternative approach would be to allow heterogeneity in all three parameters simultaneously, which would require a further expansion of the model. We hope to explore this approach in future research.

$$(4b) \quad v_t = (1 - f_t - \delta_t)^\tau v_{t-1} + \theta_t \sum_{s=1}^{\tau} (1 - f_t - \delta_t)^{s-1}$$

$$(5b) \quad H_t = f_t v_{t-1} \sum_{s=1}^{\tau} (1 - f_t - \delta_t)^{s-1} + f_t \theta_t \sum_{s=1}^{\tau} (\tau - s)(1 - f_t - \delta_t)^{s-1} + \tau \eta_t$$

where $\theta_t = \theta_t^P + \theta_t^0$ and $v_t = v_t^P + v_t^0$. By construction, $v_{t-1}^P = v_{t-1}$ and $v_{t-1}^0 = 0$.

Equations (4b) and (5b) are identical in structure to (4) and (5) from the basic model, save for the introduction of $\tau \eta_t$. The system (4b), (5b), (6) and (7) is a system of four equations in five unknowns: $f_t, \theta_t, \theta_t^0, \eta_t$, and ρ_t^0 .¹² To close the model we require an additional identifying assumption about the hires from unposted vacancies. It seems plausible that the propensity to hire from unposted vacancies varies intrinsically across establishments. For evidence in support of this, return to Figure 8, which shows the probability of a vacancy posting as a function of employment growth. The inclusion of establishment fixed effects greatly alters this function, suggesting that some establishments regularly post vacancies, while others regularly do not. Yet, accounting for establishment fixed effects in the hires and vacancy yield relations to growth does little to alter their patterns. Accordingly, we consider different propensities of unposted vacancy use via variations in ρ_t^0 . For our estimation here, however, we proceed with $\rho_t^0 = 1$, which implies that establishments with $v_{t-1} = 0$ are the only ones who hires from unposted vacancies. Figure 12 illustrates why we assume such an extreme value. As we increase ρ_t^0 , the values of f_t and θ_t also increase, but η_t decreases sharply, implying a greater share of hires from unposted vacancies. Further, $\rho_t^0 = 0.5$ yields the implausible

¹² The system (4a), (5a), (6) and (7) is an equivalent system. We use the system (4b), (5b), (6) and (7) since (4b) and (5b) closely resemble (4) and (5) and numerically we have found it easier to solve this latter system.

result that most hires (81 percent) come from unposted vacancies. Given the difficulty of formalizing the process for these types of hires (we are essentially trying to put structure on the unobserved heterogeneity of an unobserved variable), an assumption of $\rho_t^0 = 1$ proves useful because it provides us a lower-bound estimate of one of our main empirical findings: that hiring from unposted vacancies represents is a common occurrence in the labor market matching process.

4.C.2 Results for Extended Model

The results from our expanded stock-flow model are below. We again estimate the model for the aggregate time-series and for the industry and size cross-section, though the need to use estimates tabulated from microdata precludes us from using the CPS and HWI data. We tabulate estimates of H_t , v_t , and v_{t-1} , as well as v_t^0 and H_t^0 , directly from the JOLTS microdata. Remember that the latter are calculated from the subset of observations with $v_{t-1} = 0$. Our model provides estimates of θ_t and θ_t^0 which in turn yield an estimate of θ_t^P .

Time-series and cross-section results for our expanded model are in Figure 13 and Table 4, respectively. On average, vacancy flow rates for those who had a posted vacancy (2.7 percent) are considerably higher than for those who did not have a posted vacancy (1.4 percent), implying that establishments who begin the period without a vacancy posted are less likely to post one during the subsequent month. The mean job-filling rate is lower in the expanded model (3.3 percent), leading to a higher mean duration rate (30 days). Surprisingly, even though we account for the intra-month flow of vacancies, introduce heterogeneity in vacancy posting, and make the extreme assumption of $\rho_t^0 = 1$,

the monthly hires rate from unposted vacancies ($\tau \cdot \eta_t$) averages 1.2 percent of employment, representing 36 percent of all hires.

In Figure 13, vacancy flow rates for both types of firms drop slightly during the 2001 recession and gradually increases starting in mid-2003. The job-filling rate remains increasing during the 2001 recession, but is less volatile than in the basic model. Finally, hires from unposted vacancies exhibit only modest variation over our sample period, and no clear cyclical pattern emerges.

The results for the expanded model by industry and establishment size are in Table 4. The expanded model does particularly well in highlighting the differences in recruitment across these categories. For example, high turnover industries such as Resources, Construction, Retail, and Leisure and Hospitality still tend to have higher fill rates, but the disparity is considerably less. These industries also tend to have high vacancy flow rates, regardless of whether a vacancy was posted at the end of the previous month. The most notable variation across industries, though, is in the rates of hires from “unposted” vacancies. High-turnover industries have the highest rates of these hires, while industries such as Government, Health and Education, and Information, have the lowest. Similarly, Resources and Construction tend to have the highest *shares* of their hires come from unposted vacancies (59 and 58 percent, respectively), while Government and Health and Education have the lowest shares of hires of this nature (17 and 22 percent, respectively). This further reinforces the notion that different sectors use vacancies as a recruitment tool in very different ways, and that hires from unposted vacancies are an important part of the recruitment process. The evidence also suggests

that the degree to which an establishment uses a vacancy to attract workers is decreasing in the amount of turnover its industry regularly incurs.

Across size classes, the job-filling rate decreases somewhat with size, but not nearly as much as the basic model implied. Vacancy flow rates among establishments with $v_{t-1} > 0$ decreases with size in a similar manner to the vacancy flow rate in the basic model, but the flow rate among establishments with $v_{t-1} = 0$ appears to be independent of size. Hires from unposted vacancies as well as the share of hires from unposted vacancies, however, both decrease significantly with establishment size.

4.D Counterfactuals

The time-series results of both our basic and expanded models suggest that the job-filling rate, vacancy flows, and hires from unposted vacancies all play some role in accounting for the cyclical behavior of hires and vacancies. In this section, we perform some counterfactual exercises to isolate the movements in hires and vacancy rates due solely to movements in each parameter.

Our approach is straightforward. For each of the three model estimations above (basic model with JOLTS data, basic model with CPS and HWI data, and expanded model with JOLTS data), we take our parameter estimates from the previous section. We then obtain predicted estimates of v_t and h_t using equations (4) and (5) (or equations (4b) and (5b) for the expanded model) by allowing one parameter to vary over time while the other(s) remain fixed at their mean value. Thus, to get the effect of movements in the job-filling rate on v_t and h_t in the basic model, we estimate (4) and (5) using the monthly values of f_t with θ fixed at the mean of its monthly estimates. To get the effect of movements in the job-filling rate on v_t and h_t in the expanded model, we estimate (4b)

and (5b) using the monthly values of f_t with θ and η fixed at the means of their monthly estimates, and so on.

Our results are in Figures 14-16. Figure 14 depicts the results of our exercise for the basic model using the JOLTS data. The results suggest that movements in the job-filling rate account for much of the cyclical movements in both hires and vacancies. Ignoring movements in the fill rate (so that only movements in the vacancy flow rate matter) tends to overpredict the volatility of hires and underpredict the volatility of vacancies, particularly during the 2001 recession. Regressions of actual on predicted rates imply that movements in the fill rate account for 83 percent of the variation in hires and 87 percent of the variation in vacancies. Movements in the vacancy flow rate account for 69 and 75 percent, respectively.

Figure 15 presents the results of our counterfactual exercise for the basic model using the CPS and HWI data. The different data and longer time-series produce nearly identical results, with one notable difference. Movements in the job-filling rate alone continue to predict hires and vacancies well, but movements in the vacancy flow rate alone tend to substantially overpredict the cyclicalities of hires and underpredict the cyclicalities of vacancies. In fact, fixing the job-filling rate predicts an essentially acyclical vacancy rate. Regressions of actual on predicted values with these data imply that movements in the fill rate account for 46 percent of the variation in hires and 91 percent of the variation in vacancies. Movements in the vacancy flow rate account for almost none (0.2 percent) of the variation in hires and 11 percent of the variation in vacancies.

Figure 16 presents the results for the expanded model using the JOLTS data. The job-filling rate continues to predict hires and vacancies relatively well. Relative to the

basic model, the prediction power of movements in vacancy flows improves considerably. In addition, movements in hires from unposted posted vacancies also do well in predicting the variation in hires and vacancies. In regressions of actual on predicted values, movements in each account for 58-68 percent of the variation in hires, and 84-93 percent of the variation in vacancies, with movements in hires from unposted vacancies representing the upper end of each range.

Overall, the counterfactual exercises suggest a prominent role for the job-filling rate in explaining the cyclical behavior of hires and vacancies. Yet, when one accounts for the fact that many hires come from unposted vacancies, all three parameters—the fill rate, the vacancy flow rate, and the hires rate from unposted vacancies—play comparable roles in accounting for cyclical movements in hires and vacancies.

4.E Accounting for micro patterns in vacancy yield

One of the most novel aspects of studying hiring and vacancy behavior with the JOLTS data is our ability to study their patterns at the establishment-level, particularly when we relate these patterns to variations in establishment growth. In section 3, we showed highly nonlinear but increasing relationships of hires, vacancies, and the vacancy yield to growth. These relationships generally prove robust to controlling for establishment fixed effects, and in the case of the vacancy yield, cannot simply be the result of time-aggregation issues in the data. Consequently, we estimate our stock-flow model using hires and vacancy data tabulated from pooled JOLTS microdata for detailed growth rate intervals. These intervals are identical to those used to the empirical analysis. We use current and lagged observations of each establishment (where the growth rate in the current month determines which growth rate interval the observation goes into) to

obtain H_t , v_t , and v_{t-1} . We use these estimates to estimate pooled average values of f_t and θ_t for each growth rate interval. For the expanded model, we estimate v_t^0 and H_t^0 using the subset of observations with $v_{t-1} = 0$ within each interval.

The underlying conceptual model for this cross-section postulates that structural heterogeneity exists within the joint distribution of f and θ at the micro level. Specifically, suppose that firms receive systematically different draws from this joint distribution. This would create a distribution of employment growth across establishments. Using our stock-flow model, we can recover estimates of the average f and θ draws within each growth rate interval. In this respect, we are not positing a causal relationship between net growth and our model parameters, but rather an equilibrium relationship across a distribution of growth rates that emerges from a model with an underlying distributions of structural heterogeneity.

Figure 17 illustrates our estimates across establishment growth rates using the basic model. We present the daily fill rate, the monthly vacancy flow rate, and the layoff rate (obtained directly from the data and defined in the model as $\tau \cdot \delta_t$). We show the layoff rate to highlight its strong declining relation to establishment growth. Both the fill rate and the vacancy flow rate increase nonlinearly with growth, with essentially flat rates among contracting establishments that decline somewhat near zero-growth and then rise sharply for expanding establishments. The increase for the vacancy flow rate and the job-filling rate are similar, and are comparable to the empirically observed increase in the hires rate. Overall, the basic stock-flow model suggests that the observed increase in the vacancy yield with growth (Figure 9) stems from both an increase in vacancy flows as well as a sharp decline in vacancy durations.

Figure 18 illustrates the expanded model results over establishment growth rate intervals. The results are similar to the basic model in that all parameters increase nonlinearly with growth, remaining essentially constant for contracting establishments, declining for zero-growth establishments, and then increasing considerably for expanding establishments. The vacancy flow rate for establishments with $v_{t-1} > 0$ exhibits the greatest increase with growth among expanding establishments, while the vacancy flow rate for establishments with $v_{t-1} = 0$ as well as the hires rate from unposted vacancies increase at roughly the same rate.

5. Concluding Remarks

This paper makes a critical examination of the behavior of hiring and vacancies at the micro level. We use a new representative survey of establishment data, the JOLTS, supported by aggregate data sources used in previous research, to study patterns of worker recruitment at both the aggregate and micro levels. We also introduce the concept of the vacancy yield, i.e., the measure of success a vacancy has in creating a hire. We find the vacancy yield to be countercyclical, consistent with standard labor market search theory. Across industries and establishment size, we find large variations not only in vacancy yields, but also in the frequency of vacancy postings and in the propensity to hire without a posted vacancy. Finally, we find strong, nonlinear relationships between hires, vacancies, and the vacancy yield and establishment-level employment growth. Contracting establishments have low levels of all three variables, and these levels are generally constant regardless of the magnitude of employment decline. Establishments with little to no employment changes tend to have the lowest rates of hires and vacancies, as well as the lowest yields. There is a sharp discontinuity for zero-growth

establishments, while expanding establishments have relatively high levels of each variable, with all three increasing sharply with the growth rate.

While one may expect such relations for vacancies or hires, it is peculiar for the vacancy yield, since theory implies it should be independent of growth. Time aggregation no doubt plays a role, but other factors, such as variations in vacancy durations and hires from unposted vacancies may drive this finding as well. Consequently, we develop a stock-flow model of vacancy posting and worker recruitment to quantify the effects of each factor. Our model, in both its basic and expanded forms, is flexible enough to address the variations we observe over time, across industries and size-classes, and at the micro level.

In our basic model, we find that vacancy flows are considerably less volatile than vacancy stocks and that variations in the job-filling rate are countercyclical and account for much of the movements in hires, vacancies and (by construction) the vacancy yield. Our expanded model suggests that *at least* one-third of all hires come from unposted vacancies. This fraction varies widely by industry, with high-turnover industries tending to have a higher proportion of hires from unposted vacancies. Accounting for these types of hires reduces the volatility of the job-filling rate and suggests a more equitable contribution of the fill rate, the vacancy flow rate, and the hires rate from unposted vacancies in accounting for the time-series variations in hires and vacancies.

Finally, our stock-flow model suggests highly nonlinear relations of the fill rate, the vacancy flow rate and the hires rate from unposted vacancies to establishment-level growth that are comparable to the empirical relations of hires, vacancies, and the vacancy

yield to growth. This suggests that an underlying structural heterogeneity in these parameters drive much of the empirical micro-relationships.

Overall, our findings raise a variety of questions about the standard approaches to modeling recruiting and search. The findings that a large fraction of hires come from unposted vacancies and that the share of hires from posted and unposted vacancies varies in both the time-series and cross section paint a much more complex picture than what current models of labor market matching capture. Nevertheless, our findings, particularly those for the vacancy yield, which matches its theoretical implications very well, provide a framework for future work on labor market search and matching. Further, the stock-flow model we present provides a flexible platform for relating the theoretical behavior of hires and vacancies to the rich heterogeneity of their behavior observed in the microdata.

References

- Abraham, Katherine, 1983. "Structural/Frictional vs. Deficient Demand Unemployment: Some New Evidence." *American Economic Review* 73(4): 708-724.
- Abraham, Katherine, 1987. "Help wanted advertising, job vacancies, and unemployment." *Brookings Papers on Economic Activity*, 207-43.
- Blanchard, Oliver J., and Diamond, Peter, 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity*, Vol. 2: 1-60.
- Burdett, Kenneth, and Mortensen, Dale T., 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review* 39(2): 257-73.
- Coles, Melvin G., and Smith, Eric, 1996. "Cross-Section Estimation of the Matching Function: Evidence from England and Wales." *Economica* 63(2): 589-97.
- Davis, Steven J., Faberman, R. Jason, and Haltiwanger, John, 2006. "The Flow Approach to Labor Markets: New Evidence and Micro-Macro Links." *Journal of Economic Perspectives*, 20(3): 3-26.
- Davis, Steven J., John C. Haltiwanger, and Scott Schuh, *Job Creation and Destruction* (Cambridge, MA: MIT Press, 1996).
- Faberman, R. Jason, 2005. "Studying the Labor Market with the Job Openings and Labor Turnover Survey." Bureau of Labor Statistics Working Paper No. 388.
- Faberman, R. Jason, 2006. "Job Flows and the Recent Business Cycle: Not All 'Recoveries' are Created Equal." Bureau of Labor Statistics Working Paper No. 391.
- Fallick, Bruce and Charles A. Fleischman, "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows," Federal Reserve Board of Governors, mimeo (2004).
- Hall, Robert E., 2005a. "Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years," forthcoming, *2005 NBER Macroeconomics Annual*, Cambridge, MA: MIT Press.
- Hall, Robert E., 2005b. "Employment Fluctuations with Equilibrium Wage Stickiness." *American Economic Review*, 95(1): 50-65.
- Holzer, Harry J., 1994. "Job vacancy rates in the firm: An empirical analysis." *Economica* 61(1): 17-36.

Moen, Espen R., 1997. "Competitive Search Equilibrium." *Journal of Political Economy*, 105(2): 385-411.

Mortensen, Dale T., and Nagypal, Eva, 2006. "More on Unemployment and Vacancy Fluctuations." Northwestern University: mimeo.

Mortensen, Dale T., and Pissarides, Christopher A., 1994. "Job creation and job destruction and the theory of unemployment." *Review of Economic Studies* 61(3): 397-415.

Mortensen, Dale T., and Pissarides, Christopher A., 1999. "New Developments in Models of Search in the Labor Market." in Orley Ashenfelter and David Card (eds.), *Handbook of Labor Economics, Volume 3* (pp. 2567-2628). Amsterdam: Elsevier Science.

Nagypal, Eva, 2005. "Labor Market Fluctuations, On-the-Job Search, and the Acceptance Curse." Northwestern University: mimeo.

Petrongolo, Barbara, and Pissarides, Christopher A., 2001. "Looking into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature*, 39(3): 390-431.

Pissarides, Christopher A., 1985. "Short-Run Equilibrium Dynamics of Unemployment, Vacancies and Real Wages." *American Economic Review*, 75(4): 676-90.

Rogerson, Richard, Shimer, Robert, and Wright, Randall, 2006. "Search-Theoretic Models of the Labor Market: A Survey." *Journal of Economic Literature*, 43(4): 959-88.

Shimer, Robert, 2005a. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review* 95(1): 25-49.


Shimer, Robert, 2005b. "Reassessing the Ins and Outs of Unemployment." University of Chicago: mimeo.

Valetta, Robert G., 2005. "Has the U.S. Beveridge Curve Shifted Back and Why?" Federal Reserve Bank of San Francisco Working Paper No. 2005-25.

Yashiv, Eran, 2000. "The Determinants of Equilibrium Unemployment." *American Economic Review*, 90(5): 1297-1322.


Yashiv, Eran, 2006. "Search and Matching in Macroeconomics." CEPR Discussion Paper Series No. 5459.

Figure 1. Aggregate Hires, Quits, Layoffs, and Vacancy Rates, JOLTS Data




Notes: Estimates tabulated from our sample of JOLTS microdata. Rates are a percentage of average employment, as defined in the text, and are seasonally adjusted using the Census X-12 process.

Figure 2. Hires from CPS Gross Flows and Vacancies from Help Wanted Data




Notes: Hires estimates are from CPS gross flows data as tabulated by Shimer (2005, for 1967-2004 series) and Fallick and Fleischman (2004, for 1994-2004 series). Vacancies estimates come from the Help Wanted Index of the Conference Board. Shimer and HWI estimates are detrended using an HP filter with smoothing parameter of $\lambda = 10^5$. See above references and text for more details.

Figure 3. Aggregate Vacancy Yield (Hires per Vacancy), JOLTS Data




Notes: The vacancy yield is tabulated from our sample of JOLTS microdata as aggregate hires in month t divided by aggregate vacancies in posted at the end of month $t-1$.

Figure 4. Aggregate Vacancy Yield (Hires per Vacancy), CPS and HWI Data




Notes: Hires estimates are from CPS gross flows data as tabulated by Shimer (2005, for 1967-2004 series) and Fallick and Fleischman (2004, for 1994-2004 series). Vacancies estimates come from the Help Wanted Index of the Conference Board. Shimer and HWI estimates are detrended using an HP filter with smoothing parameter of $\lambda = 10^5$. The JOLTS yield is calculated using the quarterly average of the monthly hires rate. See above references and text for more details.

Figure 5. Distribution of Establishment-Level Vacancies, Employment-Weighted
(a) Vacancy Rates (Percent of Employment)




(b) Vacancy Levels (Number of Vacancies)




Note: Figures display the employment-weighted distribution of vacancy rates (upper panel) and vacancy levels (lower panel) across pooled monthly establishment observations from our JOLTS sample.

Figure 6. Hires Rate as a Function of Establishment Employment Growth




Note: The solid line represents the mean hires rate for fine intervals over the range of growth rates. The dashed line represents the mean hires rate conditional on establishment fixed effects. We derive our estimates from the pooled monthly establishment observations of our JOLTS sample. Estimates are smoothed using a centered, 5-interval moving average, with a discontinuity allowed at zero.

Figure 7. Vacancy Rate as a Function of Establishment Employment Growth




Note: The solid line represents the mean vacancy rate (measured at the end of the previous month) for fine intervals over the range of growth rates. The dashed line represents the mean vacancy rate conditional on establishment fixed effects. We derive our estimates from the pooled monthly establishment observations of our JOLTS sample. Estimates are smoothed using a centered, 5-interval moving average, with a discontinuity allowed at zero.

Figure 8. Probability of a Vacancy Posting as a Function of Establishment Employment Growth, Employment-Weighted




Note: The solid line represents the probability of a vacancy posting (measured at the end of the previous month and weighted by employment) for fine intervals over the range of growth rates. The dashed line represents the probability conditional on establishment fixed effects. We derive our estimates from the pooled monthly establishment observations of our JOLTS sample. Estimates are smoothed using a centered, 5-interval moving average, with a discontinuity allowed at zero.

Figure 9. Vacancy Yield as a Function of Establishment Employment Growth
(a) For all Observations




(b) Conditional on a Vacancy Posting




Note: In each panel, the solid line represents the number of hires in month t per vacancy posted at the end of month $t-1$ for fine intervals over the range of growth rates. The dashed line represents the number of hires per vacancy conditional on establishment fixed effects. In the upper panel, we measure the ratio as all hires in each interval divided by all vacancies in each interval, while in the lower panel, the ratio is the number of hires per vacancy only for establishments that posted at least one vacancy. We derive our estimates from the pooled monthly establishment observations of our JOLTS sample. Estimates are smoothed using a centered, 5-interval moving average, with a discontinuity allowed at zero.

Figure 10. Basic Model Monthly Parameter Estimates, JOLTS Data




Notes: Results are from our stock-flow model estimation using hires and vacancy rates tabulated from JOLTS microdata. See text for details.

Figure 11. Estimated Monthly Vacancy Fill and Flow Rates, Basic Model, CPS Gross Flow and HWI Data




Notes: Results are from our stock-flow model estimation using hires and vacancy rates tabulated from detrended estimates of CPS gross flow data (hires) and HWI data (vacancies). See text for details.

Figure 12. Expanded Model Parameter Estimates as a Function of ρ^0



Notes: Estimates are parameters from expanded model over a range of ρ^0 in the interval [0.5,1.0]. Estimates use hires and vacancy rates at the mean of their monthly values. See text for details.


Figure 13. Estimated Monthly Vacancy Fill and Flow Rates, Expanded Model




Notes: Results are from our expanded stock-flow model estimation using hires and vacancy rates tabulated from JOLTS microdata. See text for details.

Figure 14. Hires and Vacancy Rates Predicted from Fixed Parameter Values, Basic Model, JOLTS Data

(a) Actual and Predicted Hires Rate




(b) Actual and Predicted Vacancy Rate




Notes: Results are from our basic stock-flow model estimation using hires and vacancy rates tabulated from JOLTS microdata. See text for details.

Figure 15. Hires and Vacancy Rates Predicted from Fixed Parameter Values, Basic Model, CPS and HWI Data

(a) Actual and Predicted Hires Rate




(b) Actual and Predicted Vacancy Rate




Notes: Results are from our basic stock-flow model estimation using hires and vacancy rates tabulated from CPS gross flow (hires) and HWI (vacancies) data. See text for details.

Figure 16. Hires and Vacancy Rates Predicted from Fixed Parameter Values, Expanded Model, JOLTS Data

(a) Actual and Predicted Hires Rate




(b) Actual and Predicted Vacancy Rate




Notes: Results are from our expanded stock-flow model estimation using hires and vacancy rates tabulated from JOLTS microdata. See text for details.

Figure 17. Estimated Monthly Vacancy Fill and Flow Rates, Basic Model, as a Function of Establishment Employment Growth



Notes: Results are from our basic stock-flow model estimation using hires and vacancy rates tabulated from JOLTS microdata. See text for details. Estimates are smoothed using a centered, 5-interval moving average, with a discontinuity allowed at zero.

Figure 18. Estimated Monthly Vacancy Fill and Flow Rates, Expanded Model, as a Function of Establishment Employment Growth



Notes: Results are from our basic stock-flow model estimation using hires and vacancy rates tabulated from JOLTS microdata. See text for details. Estimates are smoothed using a centered, 5-interval moving average, with a discontinuity allowed at zero.

Table 1. Hires, Separations and Vacancies by Industry and Size, JOLTS Data

	h_t	s_t	v_t	h_t / v_{t-1}
<i>Nonfarm Employment</i>	3.2	3.1	2.2	1.3
<i>Major Industry</i>				
Resources	3.1	3.1	1.1	2.4
Construction	5.3	5.5	1.3	3.3
Manufacturing	2.2	2.7	1.4	1.5
Transport, Wholesale & Utilities	2.5	2.6	1.6	1.4
Retail Trade	4.3	4.2	1.9	2.0
Information	2.1	2.3	2.0	0.9
Finance, Insurance & Real Estate	2.2	2.1	2.0	1.0
Professional & Business Services	4.1	3.8	2.9	1.2
Health & Education	2.7	2.3	3.3	0.7
Leisure & Hospitality	6.1	5.9	2.8	1.9
Other Services	3.2	3.1	1.9	1.4
Government	1.5	1.2	1.8	0.8
<i>Establishment Size Class</i>				
0-9 Employees	3.1	3.2	1.4	1.6
10-49 Employees	3.9	4.0	1.9	1.8
50-249 Employees	3.8	3.7	2.2	1.5
250-999 Employees	2.9	2.8	2.4	1.1
1,000-4,999 Employees	2.0	1.9	2.7	0.7
5,000+ Employees	1.5	1.3	2.3	0.6

Notes: Estimates are tabulated from our sample of JOLTS microdata. Rates are as defined in the text.

Table 2. Distribution of Hires and Vacancies by Industry and Size, JOLTS Data

	Percent of Employment with $h_t = 0$	Percent of Employment with $v_{t-1} = 0$	Percent of Establishmen ts with $h_t = 0$	Percent of Establishmen ts with $v_{t-1} = 0$	Percent of h_t with $v_{t-1} = 0$	Percent of v_{t-1} with $h_t = 0$
<i>Nonfarm Employment</i>	35.9	46.5	81.6	87.6	42.3	18.7
<i>Major Industry</i>						
Resources	44.8	65.0	83.8	91.1	65.1	30.8
Construction	47.4	75.2	82.1	92.3	67.7	33.3
Manufacturing	34.8	46.4	73.2	84.4	43.3	15.8
Transport, Wholesale & Utilities	45.6	54.0	85.8	89.6	43.3	29.6
Retail Trade	44.0	61.0	77.6	87.7	49.7	22.4
Information	34.5	38.1	80.1	83.1	34.9	21.0
Finance, Insurance & Real Estate	45.0	50.2	87.9	90.9	41.6	22.8
Professional & Business Services	36.3	43.4	85.7	89.1	32.3	16.8
Health & Education	27.6	31.7	81.2	83.8	25.4	11.0
Leisure & Hospitality	33.6	55.6	66.4	81.8	48.6	16.6
Other Services	61.5	70.8	88.8	90.4	54.3	48.6
Government	22.4	25.7	76.8	75.7	20.6	12.9
<i>Establishment Size Class</i>						
0-9 Employees	87.6	92.3	91.7	94.7	77.5	62.8
10-49 Employees	60.9	75.2	65.9	78.7	61.1	35.5
50-249 Employees	29.3	45.7	33.1	50.0	37.3	14.1
250-999 Employees	13.3	20.6	14.5	22.4	17.7	5.2
1,000-4,999 Employees	4.7	8.3	5.5	9.5	6.9	1.4
5,000+ Employees	1.7	7.5	1.6	6.6	7.1	1.0

Notes: Estimates are tabulated from our sample of JOLTS microdata.

Table 3. Basic Model Results by Industry and Size Class, JOLTS Data

	Daily Fill Rate f_t	Monthly Flow Rate $\tau \cdot \theta_t$	Duration (Days) $1/f_t$
<i>Nonfarm Employment</i>	0.055	3.4	18.3
<i>Major Industry</i>			
Resources	0.105	3.3	9.5
Construction	0.148	6.1	6.8
Manufacturing	0.061	2.4	16.4
Transport, Wholesale & Utilities	0.059	2.7	16.9
Retail Trade	0.083	4.6	12.0
Information	0.039	2.1	25.8
Finance, Insurance & Real Estate	0.040	2.3	25.2
Professional & Business Services	0.050	4.4	19.8
Health & Education	0.029	2.7	34.3
Leisure & Hospitality	0.078	6.5	12.8
Other Services	0.058	3.4	17.3
Government	0.032	1.6	30.9
<i>Establishment Size Class</i>			
0-9 Employees	0.069	3.4	14.5
10-49 Employees	0.075	4.2	13.4
50-249 Employees	0.063	4.0	15.8
250-999 Employees	0.045	3.0	22.2
1,000-4,999 Employees	0.028	2.0	36.0
5,000+ Employees	0.025	1.6	40.1

Notes: Estimates are tabulated from our sample of JOLTS microdata.

Table 4. Expanded Model Results by Industry and Size Class, JOLTS Data

	Daily Fill Rate f_t	Flow Rate $\tau \cdot \theta_t,$ $v_{t-1} > 0$	Flow Rate $\tau \cdot \theta_t,$ $v_{t-1} = 0$	Monthly “Unposted” Hires Rate $\tau \cdot \eta_t$	Duration (Days) $1/f_t$
<i>Nonfarm Employment</i>	0.033	2.7	1.4	1.2	30.1
<i>Major Industry</i>					
Resources	0.039	2.1	0.8	1.8	25.8
Construction	0.055	4.5	1.6	3.1	18.1
Manufacturing	0.036	1.9	0.8	0.8	27.8
Transport, Wholesale & Utilities	0.036	2.2	1.1	0.9	27.6
Retail Trade	0.048	3.8	1.9	1.7	21.0
Information	0.026	1.5	1.4	0.6	37.8
Finance, Insurance & Real Estate	0.024	1.7	1.0	0.8	41.0
Professional & Business Services	0.035	4.1	1.7	1.1	28.4
Health & Education	0.022	2.4	1.2	0.6	45.5
Leisure & Hospitality	0.044	5.2	2.5	2.4	22.8
Other Services	0.030	2.5	1.4	1.4	33.4
Government	0.026	1.4	0.8	0.3	38.5
<i>Establishment Size Class</i>					
0-9 Employees	0.019	NA*	1.1	2.3	53.7
10-49 Employees	0.034	3.1	1.5	2.0	29.8
50-249 Employees	0.041	3.6	1.5	1.2	24.2
250-999 Employees	0.037	2.8	1.3	0.4	27.2
1,000-4,999 Employees	0.025	1.9	1.3	0.1	39.3
5,000+ Employees	0.023	1.4	1.6	0.1	43.4

Notes: Estimates are tabulated from our sample of JOLTS microdata. * The estimate of the vacancy flow rate for this group is imprecise and inconsistent (i.e., less than zero), mostly likely a consequence of the very small number of observations in this category, thus we suppress its result.